

Posterior Model Probabilities Computed From Model-Specific Gibbs Output

Richard J. Barker* and William A. Link†

May 27, 2011

Reversible jump Markov chain Monte Carlo (RJMCMC) extends ordinary MCMC methods for use in Bayesian multimodel inference. We show that RJMCMC can be implemented as Gibbs sampling with alternating updates of a model indicator and a vector-valued “palette” of parameters denoted ψ . Like an artist uses the palette to mix dabs of color for specific needs, we create model-specific parameters from the set available in ψ . This description not only removes some of the mystery of RJMCMC, but also provides a basis for fitting models one at a time using ordinary MCMC and computing model weights or Bayes factors by post-processing the Monte Carlo output.

*Department of Mathematics and Statistics, University of Otago, P.O. Box 56 Dunedin, New Zealand

†Patuxent Wildlife Research Center, Laurel, MD 20708, USA

We illustrate our procedure using several examples.

KEY WORDS: Reversible jump Markov chain Monte Carlo, Bayesian multimodel inference, Bayes factors, Posterior model probabilities.

1 Introduction

A natural Bayesian approach to problems of multimodel inference is to compute posterior model probabilities, or equivalently Bayes factors, given priors and data. Bayes factors involve marginal likelihoods and can be difficult to calculate. Here we address the problem of estimating posterior model probabilities, and provide a representation of reversible jump Markov chain Monte Carlo (RJMCMC) that allows us to use MCMC output obtained fitting models one at a time.

Techniques have been proposed for computing Bayes factors using MCMC output from independent chains generated for different models (for example, Chib 1995) or by using a search over the joint space of model indicators $M \in \mathcal{M}$ and model parameters $\theta_j \in \Theta_j$ given by $\mathcal{M} \times \prod_{j \in \mathcal{M}} \Theta_j$. Either approach is difficult in practice and it is common for model selection to be based instead on a deviance information criterion (DIC) (Spiegelhalter et al. 2002). However, there is no theoretical justification for using DIC to produce model weights or Bayes factors.

Reversible jump Markov chain Monte Carlo is an extension of ordinary MCMC, for multimodel inference. In this broader context, the posterior distribution under investigation describes parameters for a collection of models, rather than a single model; furthermore the posterior distribution describes model uncertainty through weights on a categorical variable we call *Model*.

A key step in implementing RJMCMC is the specification of bijections describing relationships between the parameters of various models (e.g., Green 1995; Gelman et al. 2004). RJMCMC is usually described in terms of $\binom{K}{2}$ such bijections where K is the number of models in the model set \mathcal{M} . Link and Barker (2010) outlined an alternative formulation of RJMCMC as simple Gibbs sampling, alternating between updating a palette of parameters $\boldsymbol{\psi}$, which is of the same dimension for all models, and the categorical variable *Model*. There are K bijections, one relating each model’s parameters to the palette $\boldsymbol{\psi}$; the $\binom{K}{2}$ bijections typically described are obtained from these. Careful construction of the palette and K bijections allows RJMCMC to be carried out using samples from model-specific posteriors, obtained one model at a time. Here we illustrate this approach and extend Link and Barker (2010) by showing that moves between models can be written so that they involve a direct draw from a known categorical distribution with all models in the sample space. This formulation obviates the need for use of a Metropolis-Hastings step that only allows pair-wise comparison of models and can be easier to implement than RJMCMC in its usual incarnations.

2 A description of RJMCMC

Suppose that we wish to evaluate the relative support provided by data \mathbf{y} to models M_k , $k = 1, 2, \dots, K$, these models being fully known except for parameter vectors $\boldsymbol{\theta}^{(k)}$ for which we have specified priors.

RJMCMC can be expressed as simple Gibbs sampling, with draws alternating between the categorical variable *Model* and a universal vector-valued parameter $\boldsymbol{\psi}$. We compare $\boldsymbol{\psi}$ to an artist's palette: as the artist combines colors on her palette to produce colors needed for specific applications, so components of $\boldsymbol{\psi}$ are combined to produce model-specific parameters $\boldsymbol{\theta}^{(k)}$. The important feature of RJMCMC is that the entire palette $\boldsymbol{\psi}$ is updated at each step of the Gibbs sampler (rather than simply those components relating to the present model).

Following Link and Barker (2010) the palette of parameters $\boldsymbol{\psi}$ is a vector of dimension d greater than or equal to the dimension of the most complex model in the model set. Parameter vector $\boldsymbol{\theta}^{(k)}$ can be recovered from the palette $\boldsymbol{\psi}$ by means of a known (invertible) mapping $g_k(\boldsymbol{\psi}) = \boldsymbol{\Theta}^{(k)} = (\boldsymbol{\theta}^{(k)}, \mathbf{u}^{(k)})'$. Vector $\mathbf{u}^{(k)}$ is irrelevant to model M_k , serving only to match the dimension of $\boldsymbol{\Theta}^{(k)}$ and $\boldsymbol{\psi}$, so that $g_k(\cdot)$ can be defined as a bijection. Thus if model M_2 has parameter space of dimension 7, and $d = 10$, vector $\mathbf{u}^{(2)}$ will have dimension 3.

Note that the $\binom{K}{2}$ bijections typically required for RJMCMC are induced by our K bijections between the palette and model-specific parameter spaces:

for models M_j and M_k we have $g_{jk}(\Theta^{(j)}) \equiv g_k \circ g_j^{-1}(\Theta^{(j)}) = g_k(\boldsymbol{\psi}) = \Theta^{(k)}$.

We must specify a prior for $\boldsymbol{\psi}$ and in doing so accommodate model specific priors $[\boldsymbol{\theta}^{(k)} | Model = M_k]$, which for simplicity we write as $[\boldsymbol{\theta}^{(k)} | M_k]$. We have

$$\begin{aligned} [\boldsymbol{\Theta}^{(k)} | M_k] &= [\boldsymbol{\theta}^{(k)}, \mathbf{u}^{(k)} | M_k] \\ &= [\boldsymbol{\theta}^{(k)} | M_k] [\mathbf{u}^{(k)} | \boldsymbol{\theta}^{(k)}, M_k]. \end{aligned}$$

All that is needed is a specification of $[\mathbf{u}^{(k)} | \boldsymbol{\theta}^{(k)}, M_k]$; given that $\mathbf{u}^{(k)}$ has no role in inference, it will be convenient to assume it is conditionally independent of $\boldsymbol{\theta}^{(k)}$, so that $[\mathbf{u}^{(k)} | \boldsymbol{\theta}^{(k)}, M_k] = [\mathbf{u}^{(k)} | M_k]$. The specific choice does not matter, except for tuning the RJMCMC algorithm. From $[\boldsymbol{\Theta}^{(k)} | M_k]$ we obtain $[\boldsymbol{\psi} | M_k]$ using the change of variables theorem in terms of a prior $f_k(\boldsymbol{\Theta}^{(k)}) = [\boldsymbol{\theta}^{(k)}, \mathbf{u}^{(k)} | M_k]$. The prior on $\boldsymbol{\psi}$ is then

$$[\boldsymbol{\psi}] = \sum_k [\boldsymbol{\psi} | M_k] [M_k]$$

where

$$[\boldsymbol{\psi} | M_k] = f_k(g_k(\boldsymbol{\psi})) \left| \frac{\partial g_k(\boldsymbol{\psi})}{\partial \boldsymbol{\psi}} \right|. \quad (1)$$

Under this formulation, Gibbs sampling consists of cyclical sampling of full conditional distributions, alternating between draws from $[\boldsymbol{\psi} | M_k, \mathbf{y}]$ to update $\boldsymbol{\psi}$, and from $[\{M_1, \dots, M_K\} | \boldsymbol{\psi}, \mathbf{y}]$ to update *Model*.

Updating ψ

To draw from the full conditional $[\psi | M_k, \mathbf{y}]$ it suffices to draw from $[\boldsymbol{\theta}^{(k)} | M_k, \mathbf{y}]$ and $[\mathbf{u}^{(k)} | M_k]$ and then to apply the inverse transformation $g_k^{-1}(\boldsymbol{\Theta}^{(k)})$ to obtain a draw for ψ .

The draw from $[\boldsymbol{\theta}^{(k)} | M_k, \mathbf{y}]$ can either be made directly, if the distribution is of convenient form, or by simulation. Another possibility, often an attractive alternative, is to take a random draw from the stored MCMC output of an earlier analysis of model M_k .

Updating the model

The full-conditional for *Model* is categorical with probabilities:

$$\Pr(M_k | \cdot) = \frac{[\mathbf{y} | \psi, M_k][\psi | M_k][M_k]}{\sum_j [\mathbf{y} | \psi, M_j][\psi | M_j][M_j]} ,$$

for $k = 1, \dots, K$. If we are willing to calculate all of these probabilities, we can update *Model* by a direct draw from this full-conditional distribution. Chain means of model indicators $I(\text{Model} = M_k)$ converge to the full conditional model probabilities, but greater efficiency is available by using chain means of the full conditional model probabilities, which also converge to the posterior model probabilities.

As an alternative, we can update model indicators by a Metropolis-Hastings step if the model candidate generator only allows limited transitions, for example, to a near neighbor in a graphical model sense. The advantage

of this approach is that we compute a smaller set of categorical probabilities, corresponding to the neighborhood set. In this case we must compute posterior model probabilities by chain means of indicators $I(\text{Model} = M_k)$.

Expressing RJMCMC as simple Gibbs sampling provides the key innovation of our formulation: it allows us to fit models one at a time using ordinary MCMC and then compute model weights or Bayes factors by post-processing the Monte Carlo output. Thus, we have a simple 2-stage procedure that can be used for computing model probabilities:

Stage 1: Produce samples of $[\boldsymbol{\psi}|\mathbf{y}, M_k]$ for each k .

Begin by sampling $[\boldsymbol{\theta}^{(k)}|\mathbf{y}, M_k]$. This can be accomplished by running an MCMC sampler for M_k , processing it in the usual way, discarding any initial burn-in iterations, and storing the results. (In cases where the posterior distribution for $\boldsymbol{\theta}^{(k)}$ is of a known and easily sampled form, we do so.) For each sampled value $\boldsymbol{\theta}^{(k)}$, independently sample an auxiliary variable $\mathbf{u}^{(k)}$ from $[\mathbf{u}^{(k)}|M_k]$ and calculate $\boldsymbol{\psi} = g_k^{-1}((\boldsymbol{\theta}^{(k)}, \mathbf{u}^{(k)})')$. The collection of sampled values $\boldsymbol{\psi}$ is a sample of $[\boldsymbol{\psi}|\mathbf{y}, M_k]$.

Stage 2: Post-process the model specific outputs.

Posterior model probabilities can be computed in one of two ways. The first method is based on generating a Markov chain of the categorical variable

$Model$ that can be used as a posterior sample from $[Model|\mathbf{y}]$. The second method is based on generating a Markov chain of between-model transition probabilities that can be used to estimate the between-model transition matrix. The steady state marginal distribution of this matrix corresponds to the posterior model probabilities.

Method 1 A posterior sample $Model^{(1)}, Model^{(2)}, \dots, Model^{(j)}, \dots$ can be generated as follows:

- (a) Initialize $Model$, say with $Model^{(1)} = M_1$.
- (b) Iterate from $j = 1$ to some large number J , and at each step:
 - (i) If $Model^{(j)} = M_k$, draw a value $\boldsymbol{\psi}^{(j)}$ from the stored sample of $[\boldsymbol{\psi}|\mathbf{y}, M_k]$ from Stage 1.
 - (ii) Compute $\pi_k^{(j)} = \Pr(Model = M_k|\boldsymbol{\psi}^{(j)}, \mathbf{y})$, for each k . This calculation requires the Jacobian of the transformation $\boldsymbol{\Theta}^{(k)} = g_k(\boldsymbol{\psi})$ as in eq. (1), evaluated at $\boldsymbol{\psi}^{(j)}$.
 - (iii) Sample $Model^{(j+1)}$ from a categorical distribution with sample space $\{M_1, M_2, \dots, M_K\}$ and probability vector $\boldsymbol{\pi}^{(j)} = (\pi_1^{(j)}, \pi_2^{(j)}, \dots, \pi_K^{(j)})'$.

The relative frequency with which $Model^{(j)} = M_k$ approximates the posterior model probability for M_k . A better estimate (Rao-Blackwellized) is the chain mean of values $\pi_k^{(j)}$.

Method 2 A further improvement on this approximation can be made by marginalizing at stage 1, obviating the need for the construction of a Markov chain of model indicators in the second stage. For each model we can compute $\Pr(M_k|\boldsymbol{\psi}, M_h)$ forming a Markov chain of transition probabilities from model h to model k ($h, k \in 1, \dots, K$). These can then be averaged to form an approximation to the stochastic matrix governing model-to-model transitions. Given an estimate of this transition matrix we can obtain corresponding estimates of the posterior model probabilities as the limiting distribution obtainable by normalizing the left eigenvector of the transition matrix associated with the eigenvalue 1.0 (Seber 2008).

3 Examples

3.1 Radiata pine data

Carlin and Chib (1995), Han and Carlin (2001), and many others analyze data taken from Williams (1959). The response variable y_i is the maximum compressive strength parallel to the grain for 42 radiata pine boards. Two explanatory variables are considered: the first is the specimen's density, x_i , and the second is the specimen's density having adjusted for resin content, z_i . Resin increases the density of boards without increasing their compressive

strength. Carlin and Chib (1995) considered two models:

$$\text{Model 1: } y_i = \alpha + \beta(x_i - \bar{x}) + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma_x^2)$$

and

$$\text{Model 2: } y_i = \gamma + \delta(z_i - \bar{z}) + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma_z^2).$$

In both cases the errors ε are assumed iid among observations, conditional on the parameters.

As priors, Carlin and Chib (1995) used $N((3000, 185)', \text{diag}(10^6, 10^4))$ priors on $(\alpha, \beta)'$ and $(\gamma, \delta)'$, and inverse gamma priors on σ_x^2 and σ_z^2 , both having mean and standard deviation equal to 300^2 . This quirky choice of priors was made to be vague but with expectations corresponding roughly to the parameter estimates obtained by fitting the model by least squares.

We fitted each of these models independently using BUGS (Lunn et al. 2000) and the above priors, running three chains of 60,000 each with distinct starting values. Discarding the first 10,000 of each chain as a burn-in left us with a posterior sample of 150,000 for $\boldsymbol{\theta}^{(1)}$ and $\boldsymbol{\theta}^{(2)}$. We coded a reversible jump algorithm in which $g_1(\boldsymbol{\psi}) = (\alpha, \beta, \sigma_x^2)'$ and $g_2(\boldsymbol{\psi}) = (\gamma, \delta, \sigma_z^2)'$. In this case, $[\boldsymbol{\psi}|M_1] = [\boldsymbol{\psi}|M_2]$, and the model update is based on the relative values

of the likelihoods weighted by the model priors $\Pr(M_k)$:

$$\Pr(M_k|\mathbf{y}, \boldsymbol{\psi}) = \frac{\Pr(M_k)e^{\frac{-1}{2\psi_3} \sum_{i=1}^{42} (y_i - \mu_{ik})^2}}{\Pr(M_1)e^{\frac{-1}{2\psi_3} \sum_{i=1}^{42} (y_i - \mu_{i1})^2} + \Pr(M_2)e^{\frac{-1}{2\psi_3} \sum_{i=1}^{42} (y_i - \mu_{i2})^2}},$$

where

$$\mu_{ik} = \begin{cases} \psi_1 + \psi_2(x_i - \bar{x}) & k = 1 \\ \psi_1 + \psi_2(z_i - \bar{z}) & k = 2 \end{cases}.$$

Following Han and Carlin (2001) we assigned model priors of $\Pr(M_1) = 0.9995$ and $\Pr(M_2) = 0.0005$ to ensure that the two models were visited in roughly equal proportion. Starting at model 1 or model 2 the chain for the posterior model probability converges rapidly (Figure 1). After running the two chains for 200,000 iterations and discarding the first 100,000 as a burn-in, our estimate of the posterior model probability was 0.709 corresponding to a Bayes factor BF_{21} of 4870. These are in close agreement with the exact values of $\Pr(M_2|\mathbf{y}) = 0.70865$ and $BF_{21} = 4862$ reported by Han and Carlin (2001).

Figure 1 about here

Using our second method, we sampled 200,000 values of $\boldsymbol{\psi}$ from each chain h , and for the i th sample we calculated $\Pr(M_k|\boldsymbol{\psi}^{(i)}, M_h)$ ($k = 1, 2$). Averaging across i we obtain an estimated transition matrix of:

$$\begin{pmatrix} 0.6003 & 0.3997 \\ 0.1651 & 0.8349 \end{pmatrix}$$

with steady-state marginal distribution of $(0.2924, 0.7076)'$, corresponding to $BF_{21} = 4838$.

3.2 Trout return rates

Link and Barker (2006) report an analysis based on fitting logistic regression models to the return rates for brown trout expressed in terms of sex S_i and length L_i effects. Modeling the return indicator $y_i \sim \text{Bern}(p_i)$ they considered five models:

1. $\eta_i = \beta_0$
2. $\eta_i = \beta_0 + \beta_1 S_i$
3. $\eta_i = \beta_0 + \beta_2 L_i$
4. $\eta_i = \beta_0 + \beta_1 S_i + \beta_2 L_i$
5. $\eta_i = \beta_0 + \beta_1 S_i + \beta_2 L_i + \beta_3 S_i L_i$

where $\eta_i = \text{logit}(p_i)$.

Link and Barker (2006) used the following priors on parameters:

$$[\boldsymbol{\beta}_k | V, M_k] = \begin{cases} N(0, V^{-1}) & k = 1 \\ N(0, (2V)^{-1}) & k = 2 \\ N(0, (2V)^{-1}) & k = 3 \\ N(0, (3V)^{-1}) & k = 4 \\ N(0, (4V)^{-1}) & k = 5 \end{cases}$$

where V has a $Ga(3.29, 7.80)$ prior distribution. This choice was motivated by the observation that if $\text{logit}(p) \sim N(0, V^{-1})$ and $V \sim \Gamma(3.29, 7.80)$, then the marginal distribution of p is very nearly uniform on $[0,1]$. With S_i and L_i having been standardized, these choices of priors ensure that the prior on $e^{n_i}/(1 - e^{n_i}) = p_i$ is approximately $U(0, 1)$ for $S_i = \pm 1$ and $L_i \pm 1$.

Palette and bijections

Each element of $\boldsymbol{\psi}$ is directly associated with either an element of the beta vector or with a supplemental variable u (Table 1): The parameter V is part

	Model				
$\boldsymbol{\psi}$	1	2	3	4	5
ψ_1	β_0	β_0	β_0	β_0	β_0
ψ_2	u_1	β_1	u_1	β_1	β_1
ψ_3	u_2	u_1	β_2	β_2	β_2
ψ_4	u_3	u_2	u_2	u_1	β_{12}

Table 1: Association between elements of $\boldsymbol{\psi}$ and elements of $\boldsymbol{\beta}_k$, specific parameters for model M_k and supplemental variables \boldsymbol{u}_k used in model M_k for matching the parameter dimension to $\boldsymbol{\psi}$.

of the prior specification and is common to all models so we chose to leave it out of the palette specification, although this is not necessary. Updates for V were stored when each model was fitted.

For a particular model, the priors on the supplemental variables were the same as the priors used for the β coefficients in that model, and in each case the Jacobian of the transformation from $\boldsymbol{\theta}^{(k)}$ to $\boldsymbol{\psi}$ is an identity matrix of dimension 5.

As an example, Model 1 (constant only) has parameter vector $\boldsymbol{\theta}^{(1)} = \beta_0$ with supplemental variables $\mathbf{u}^{(1)} = (u_1, u_2, u_3)'$. Thus

$$g_1(\boldsymbol{\psi}) = g_1 \begin{pmatrix} \psi_1 \\ \vdots \\ \psi_4 \end{pmatrix} = \boldsymbol{\Theta}^{(1)} = \begin{pmatrix} \beta_0 \\ u_1 \\ u_2 \\ u_3 \end{pmatrix} = \begin{pmatrix} \psi_1 \\ \vdots \\ \psi_4 \end{pmatrix} ,$$

leading to:

$$\begin{aligned} [\boldsymbol{\psi}|M_1, V] &= f_1(g_1(\boldsymbol{\psi})) \left| \frac{\partial g_1(\boldsymbol{\psi})}{\partial \boldsymbol{\psi}} \right| \\ &= N(\psi_1; 0, V^{-1}) \times N(\psi_2; 0, V^{-1}) \times N(\psi_3; 0, V^{-1}) \times N(\psi_4; 0, V^{-1}) \\ &\quad \times Ga(V; 3.29, 7.80) \times |\mathbf{I}_5| . \end{aligned}$$

Repeating this process for each model we obtain the model-specific priors:

$$[\boldsymbol{\psi}|M_k, V] = \prod_{i=1}^4 N(\psi_i; 0, (n_k V)^{-1}) \times Ga(V; 3.29, 7.80) \times |\mathbf{I}_5|$$

where n_k is the dimension of the vector $\boldsymbol{\beta}^{(k)}$.

For generating a chain of model indicators, we used a direct draw from the full conditional:

$$\Pr(M_k|\boldsymbol{\psi}, V) = \frac{\Pr(M_k) \prod_{i=1}^4 \sqrt{\frac{n_k V}{2\pi}} e^{\frac{-n_k V}{2} \psi_i^2} \prod_{j=1}^{1961} p_j^{(k)} (1 - p_j^{(k)})}{\sum_{h=1}^5 \Pr(M_h) \prod_{i=1}^4 \sqrt{\frac{n_h V}{2\pi}} e^{\frac{-n_h V}{2} \psi_i^2} \prod_{j=1}^{1961} p_j^{(h)} (1 - p_j^{(h)})}$$

where $\text{logit}(p_i^{(k)}) = x_i' \boldsymbol{\beta}^{(k)}$.

To estimate the Bayes factors we first fitted the five models, in each case combining results from three different chains of length 500,000 after discarding a burn-in. We then generated five chains using our Gibbs sampler for the model indicator, starting each chain with a different model. Following Link and Barker (2006) we first tuned the Gibbs sampler to visit each model in roughly equal proportion. Mixing of the model indicators appears rapid (Figure 2) and agreement with the Link and Barker (2006) estimates is good after combining the results from the second half of 200,000 iterations of the five chains (Table 2).

j	BF_{1j}	$\text{Pr}(M_j \mathbf{y})$
1	1 (1)	0.893 (0.894)
2	31.3 (31.7)	0.029 (0.028)
3	12.3 (12.4)	0.073 (0.072)
4	274.6 (281.7)	0.003 (0.003)
5	383.4 (390.1)	0.002 (0.002)

Table 2: Estimates of Bayes factors BF_{1j} for comparing models 1 and j and estimates of posterior model probabilities under constant prior model probabilities $\text{Pr}(M_j) = 0.2$. Corresponding estimates from Link and Barker (2006) are given in parentheses.

Figure 2 about here

For method two we drew a sample of 10,000 values for $\boldsymbol{\psi}$ from each chain

¹ leading to an estimate of the transition matrix of:

$$\begin{pmatrix} 0.8172 & 0.0870 & 0.0847 & 0.0088 & 0.0024 \\ 0.0858 & 0.8086 & 0.0107 & 0.0755 & 0.0195 \\ 0.0854 & 0.0102 & 0.8233 & 0.0759 & 0.0052 \\ 0.0081 & 0.0749 & 0.0781 & 0.7884 & 0.0504 \\ 0.0026 & 0.0176 & 0.0057 & 0.0498 & 0.9244 \end{pmatrix}$$

with steady-state marginal distribution

$$\begin{pmatrix} 0.1986 \\ 0.1975 \\ 0.2016 \\ 0.1989 \\ 0.2034 \end{pmatrix}.$$

3.3 Simple binomial

In both of the above examples, the bijections from $\boldsymbol{\psi}$ to $\boldsymbol{\theta}$ are simple 1-1 mappings with the Jacobian of the transformation an identity matrix. Now consider an example where $Y_i \sim B(N_i, p_i)$ and we have observations $y_1 = 8$, $n_1 = 20$, $y_2 = 16$, and $n_2 = 30$. What is the evidence for $p_1 \neq p_2$ against $p_1 = p_2 = \pi$? To compute an appropriate Bayes factor we fit two models:

1. Model 1: (p_1, p_2) with independent $Be(\alpha_p, \beta_p)$ priors

¹Only 10,000 samples were drawn due to the large RAM requirements on the desktop. This number can easily be increased by writing batches of such draws to the hard-drive.

2. Model 2: $p_1 = p_2 = \pi$ with a $Be(\alpha_\pi, \beta_\pi)$ prior on π .

For model 1, we assign $\boldsymbol{\psi} = (p_1, p_2)'$. It seems natural in moving from model 1 to model 2 that the average $\bar{\psi} = (\psi_1 + \psi_2)/2$ should provide a good candidate for π . Thus, our bijections can be written as:

$$\text{Model 1: } \mathbf{I}_2 \times \boldsymbol{\psi} = \begin{pmatrix} p_1 \\ p_2 \end{pmatrix}$$

and

$$\text{Model 2: } \begin{pmatrix} 1/2 & 1/2 \\ 0 & 1 \end{pmatrix} \times \begin{pmatrix} \psi_1 \\ \psi_2 \end{pmatrix} = \begin{pmatrix} \pi \\ u \end{pmatrix},$$

where \mathbf{I}_2 is a 2×2 identity matrix and u an appropriate supplemental variable

Our Gibbs sampler then proceeds as follows:

1. Within models the full conditional distributions for model-specific parameters are of known form since we have conditional (on the model) conjugacy:

- Under Model 1 we sample $p_1 \sim Be(8 + \alpha_p, 12 + \beta_p)$ and $p_2 \sim Be(16 + \alpha_p, 14 + \beta_p)$ and then compute $\boldsymbol{\psi} = (p_1, p_2)'$.
- Under Model 2 we sample $\pi \sim Be(24 + \alpha_\pi, 26 + \beta_\pi)$ and $u \sim Be(\alpha_u, \beta_u)$ and then set $\psi_1 = 2\pi - u$ and $\psi_2 = u$.

2. Between models we set:

$$\begin{aligned}
& - \Pr(M_1|\cdot) \propto \Pr(M_1) \times \psi_1^{8+\alpha_p} (1-\psi_1)^{12+\beta_p} \psi_2^{16+\alpha_p} (1-\psi_2)^{14+\beta_p} \mathbf{1}_{\psi_1 \in (0,1)} \mathbf{1}_{\psi_2 \in (0,1)} \\
& - \Pr(M_2|\cdot) \propto \Pr(M_2) \times \bar{\psi}^{24+\alpha_\pi} (1-\bar{\psi})^{26+\beta_\pi} \mathbf{1}_{\bar{\psi} \in (0,1)} \mathbf{1}_{\psi_2 \in (0,1)} \times \frac{1}{2}
\end{aligned}$$

where $\mathbf{1}_E$ denotes the indicator of the event E and the proportionality constant is the same for each model. We then sample the model indicator by a direct draw from a categorical distribution with sample space $\{1, 2\}$ and parameter vector $(\pi_1, \pi_2)'$ where

$$\pi_j = \frac{\Pr(M_j|\cdot)}{\Pr(M_1|\cdot) + \Pr(M_2|\cdot)}.$$

To fit the models we used as prior parameters $\alpha_p = \beta_p = \alpha_\pi = \beta_\pi = 1$ (i.e., independent $U(0, 1)$ priors). We also set $\alpha_u = \beta_u = 15$ so that draws for u were similar to draws for p_2 . Convergence of the chain for the posterior probability of model 2 was rapid (Figure 3). Combining results from 100,000 iterations of the two chains we obtained $\hat{B}_{21} = 1.92$ and $\Pr(M_2) = 0.658$. For both models the marginal distribution of the data is straight-forward to compute and the exact solution for the posterior model probability is 0.6580.

Figure 3 about here

Using method two with a sample of 100,000 values of ψ from each chain we estimate the transition matrix as:

$$\begin{pmatrix} 0.4318 & 0.5682 \\ 0.2951 & 0.7049 \end{pmatrix}$$

with steady-state marginal distribution $(0.34230.6577)'$.

4 Discussion

Bayesian inference offers an appealing framework for multimodel inference but the difficulties of computing Bayes factors, or equivalently posterior model probabilities, can be a barrier to implementation. Being able to independently fit models and then post-process them using RJMCMC as we have described here offers a partial solution to the problem.

An issue often raised in objection to Bayesian multimodel inference (BMI) based on Bayes factors is that one must assume that the true model is in the model set. As we have argued elsewhere (Link and Barker 2006, 2010), this is a red-herring – conditioning on a model set is no less innocuous than conditioning on a model as must be done for any form of statistical inference. Conditioning on models and model sets is done for operational convenience - we no more believe that truth is in our model set than we believe that the model $y_i \stackrel{iid}{\sim} N(\mu, \sigma^2)$ can ever be a true and complete representation of any set of data.

A more serious issue with BMI is priors on parameters; it is well-known that Bayes factors are sensitive to choice of priors, particularly vague priors. Our view is that priors should be chosen so that common features of interest in each model have the same prior uncertainty associated with them. An attempt at such an approach is illustrated by the West Coast trout example

in which priors were constructed based on the logit of the return probability for trout that had typical values of the covariates. Such an approach we have previously referred to as “nonpreferential” (Link and Barker 2008).

Choice of efficient bijections for moving between models requires some thought. Although our approach simplifies this problem to one of choosing K such bijections, choices must be made. Features that are of interest and common across models can be exploited in choosing bijections as well as providing a basis for constructing non-preferential priors. Generalized linear model formulations such as represented in our trout example offers one means for constructing bijections. One area of possible fruitful investigation in this context is that our palette representation of RJMCMC appears to be connected to the use of importance link functions MacEachern and Peruggia (2000). There may be benefits from considering this connection from the point of view of determining transformations $g(\psi)$ in our representation that lead to more efficient Monte Carlo estimation of posterior model probabilities.

Our description of RJMCMC as simple Gibbs sampling with a direct draw from a known distribution for model probabilities is a further useful simplification. Moves can be made to any model in the set \mathcal{M} using samples from the full-conditional distribution for model indicators; we are not restricted to moves between pairs of models. Methods that involve moves to neighbours have been used to automate search across very high dimensional model space. We are skeptical about the value of such algorithms as they induce a particular prior on parameters. Such default constructions may lead to priors that

are prejudicial in which case posterior model probabilities would be more a reflection of these prior prejudices than data-informed posterior weighting.

References

- Carlin, B. P. and Chib, S. (1995), “Bayesian model choice via Markov chain Monte Carlo methods.” *Journal of the Royal Statistical Society, Series B.*, 57, 473–484.
- Chib, S. (1995), “Marginal likelihood from the Gibbs output.” *Journal of the American Statistical Association*, 90, 1313–1321.
- Gelman, A. B., Carlin, J. S., Stern, H. S., and Rubin, D. B. (2004), *Bayesian data analysis. Second edition.*, Chapman and Hall/CRC, Boca Raton.
- Green, P. J. (1995), “Reversible jump Markov chain Monte Carlo computation and Bayesian model determination.” *Biometrika*, 82, 711–732.
- Han, C. and Carlin, B. (2001), “Markov chain Monte Carlo methods for computing Bayes factors: a comparative review.” *Journal of the American Statistical Association*, 96, 1122–1133.
- Link, W. and Barker, R. (2008), “Bayes factors and multimodel inference,” *Environmental and Ecological Statistics*, 3, 597–618.
- Link, W. A. and Barker, R. J. (2006), “Model weights and the foundations of multi-model inference,” *Ecology*, 87, 2626–2635.

- (2010), *Bayesian Inference with ecological applications*, Academic Press.
- Lunn, D., Thomas, A., Best, N., and Spiegelhalter, D. (2000), “WinBUGS – a Bayesian modelling framework: concepts, structure, and extensibility.” *Statistics and Computing*, 10, 325–337.
- MacEachern, S. N. and Peruggia, M. (2000), “Importance link function estimation for Markov chain Monte Carlo methods,” *Journal of Computational and*, 9, 99–121.
- Seber, G. A. F. (2008), *A matrix handbook for statisticians*, John Wiley and Sons, Hoboken, New Jersey.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002), “Bayesian measures of model complexity and fit.” *Journal of the Royal Statistical Society series B*, 64, 583–639.
- Williams, E. (1959), *Regression Analysis.*, Wiley, New York.




Figure 1: Index plot of the cumulative posterior probability $p = \Pr(M = 2)$ starting with model 1 (red) or model (2). The horizontal black line corresponds to the exact result.

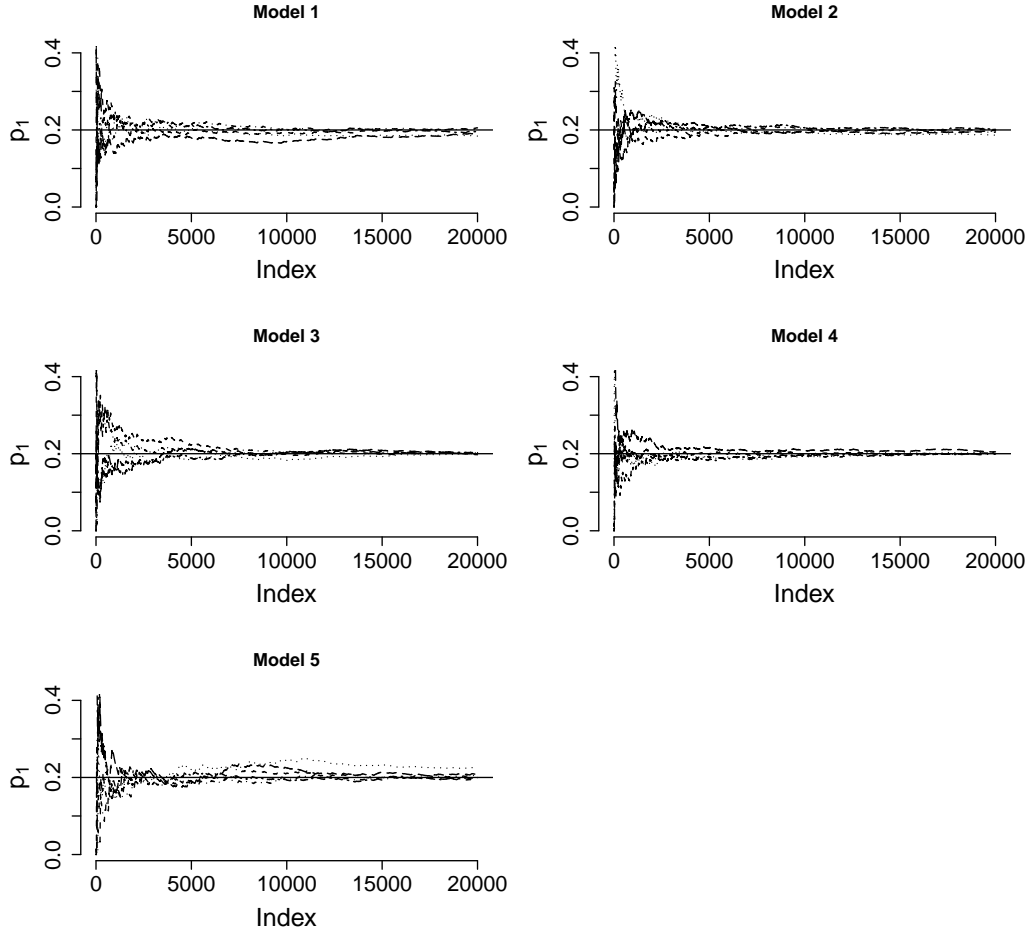


Figure 2: Index plot of the cumulative posterior model probabilities. Each plot represents a different model probability and the different colored chains represent different starting values. The black line corresponds to the value targeted during tuning.




Figure 3: Index plot of the cumulative posterior probability $p = \Pr(M = 2)$ starting with model 1 (red) or model (2). The horizontal black line corresponds to the exact result.